

CLAIMS

What is claimed is:

1. A method comprising:

determining that a pre-registered remote direct memory access (RDMA) buffer

5 has insufficient size to transfer data;

provisioning a larger RDMA buffer; and

transferring the data to a network using the larger RDMA buffer.

2. The method of claim 1, further comprising:

sending a control message indicating that a receiver is to provision a larger

10 RDMA buffer to receive the data; and

prior to said transferring, receiving an acknowledgement message indicating that the receiver has provisioned the larger RDMA buffer.

3. The method of claim 1, wherein said determining comprises comparing a size of the data to a predetermined threshold.

15 4. The method of claim 3, further comprising comparing sizes of a plurality of elements of an input-output vector to the predetermined threshold.

5. The method of claim 1, wherein said provisioning comprises allocating and registering the larger RDMA buffer during a communication phase.

6. The method of claim 1, wherein said provisioning comprises:

20 unregistering the pre-registered RDMA buffer; and

freeing the memory used by the pre-registered RDMA buffer.

7. The method of claim 1, wherein said transferring comprises:

copying data from a source to the larger RDMA buffer; and

performing an RDMA transfer from the larger RDMA buffer to the network.

25 8. An article of manufacture comprising:

a machine-accessible medium that provides instructions that if executed result in a machine performing operations including,

determining that a size of data to be transferred by remote direct memory access (RDMA) is larger than a predetermined threshold; and

30 allocating and registering an RDMA buffer having a size larger than the predetermined threshold during a communication phase of an RDMA data transfer.

9. The article of manufacture of claim 8, wherein the machine-accessible medium further provides instructions that if executed result in the machine performing

operations comprising:

sending a message indicating that a receiver is to provision an RDMA buffer having a size larger than the predetermined threshold; and

5 receiving an acknowledgement message indicating that the receiver has provisioned the larger RDMA buffer.

10. The article of manufacture of claim 8, wherein the machine-accessible medium further provides instructions that if executed result in the machine performing operations including comprising:

copying data from a source to the RDMA buffer; and

10 performing an RDMA transfer from the RDMA buffer to a receiving node.

11. A system comprising:

an interconnect;

one or more processors coupled with the interconnect;

15 a dynamic random access memory (DRAM) coupled with the interconnect to store data;

a network interface device coupled with the interconnect to transfer data to a network by using an Ethernet protocol;

machine-readable instructions stored in the DRAM that if executed result in a machine performing operations comprising:

20 determining that a pre-registered remote direct memory access (RDMA) buffer is too small to transfer data;

provisioning an RDMA buffer having a size larger than the pre-registered RDMA buffer; and

transferring the data to a network using the provisioned RDMA buffer.

25 12. The method of claim 11, further comprising:

sending a message indicating that a receiver is to provision an RDMA buffer having a size larger than the pre-registered RDMA buffer to receive the data; and

prior to said transferring, receiving an acknowledgement message indicating that the receiver has provisioned the RDMA receive buffer.

30 13. The method of claim 11, wherein said determining comprises comparing a size of the data to a predetermined threshold, wherein said provisioning comprises allocating and registering the larger RDMA send buffer during a communication phase, and wherein said transferring comprises copying data from a source to the larger

RDMA send buffer.

14. The method of claim 11, wherein the pre-registered RDMA buffer has a size ranging from 100 to 2,000 bytes, and wherein the provisioned RDMA buffer has a size ranging from 1,000 to 200,000 bytes.

5 15. A method comprising:

receiving a control message indicating to provision a remote direct memory access (RDMA) buffer having a size larger than a pre-registered RDMA buffer;
provisioning the RDMA buffer;
receiving data into the provisioned RDMA buffer.

10 16. The method of claim 15, further comprising sending an acknowledgement message indicating that the RDMA buffer has been provisioned.

17. The method of claim 15, wherein said provisioning comprises allocating and registering the RDMA buffer during a communication phase.

15 18. The method of claim 15, wherein said provisioning comprises:
unregistering the pre-registered RDMA buffer; and
freeing the memory used by the pre-registered RDMA buffer.

19. The method of claim 15, further comprising copying the data from the provisioned RDMA buffer to a destination.

20 20. The method of claim 15, implemented in a network device comprising:
a bus;
a dynamic random access memory coupled with the bus to store data; and
a processor having multiple cores.